MILITARY REQUIREMENTS FOR PACKET-SWITCHED NETWORKS AND THEIR IMPLICATIONS FOR PROTOCOL STANDARDIZATION

Dr. Vinton G. Cerf
Defense Advanced Research Projects Agency

and

Dr. Robert E. Lyons
DCA/Defense Communications Engineering Center

ABSTRACT

This paper outlines the nature of the military requirements for packet-switched data communications and contrasts these requirements with those of commercial, industrial and private users served by common carrier public packet-switched data networks. Current US Department of Defense policy on the use of commercial standards in defense systems is reviewed. Deficiencies in existing standards for military applications are identified and implications for the development of a suitable military data communications architecture are drawn. A strategy of military standardization is outlined which adopts commercial standards where they apply and supplements these with special military standards as needed to satisfy unique military requirements. The thesis of this paper is illustrated by a discussion of the applicability of the CCITT X.25 protocol standard to military communications.

1. INTRODUCTION

Military requirements for packet-switched networks include many of the requirements for public networks and others besides. It is interesting that some of the research in the early sixties, which led to the original development of the packet-switching concept of communications, was motivated by the special requirements of the military. However, packet-switching systems, as they have evolved into networks for providing public data communications services, have not always paid attention to military requirements. It is not surprising, therefore, that some of the standards which have been adopted for public data networks provide inadequate service for the military user. This fact creates a dilemma for those concerned with the deployment of packet-switched communications for the military: On the one hand, those responsible for military communications systems would like to use prevailing standards because they would therefore be able to capitalize upon hardware, software and service offerings which are readily available at reasonable cost, and would be interoperable with the rest of the world of packet-switched communications, thereby providing ubiquitous means for service restoral in case of emergency. On the other hand, if the adoption of the standards of public data networks entails too much of a sacrifice of the military's ability to perform its mission, then the utility of public data network standards is low despite potential cost advantages.

The purpose of this paper is to explore this dilemma in the context of the policy presently being followed by the US Department of Defense with respect to adherence to public packet-switched communications standards. To provide context, the paper begins with a summary of the history of packet-switching technology. Next the paper briefly reviews data communications requirements with emphasis on those which are uniquely military. The DoD policy on the use of commercial standards is reviewed. As a specific example, the applicability of the International Consultative Committee on Telegraphy and Telephony (CCITT) X.25 protocol to military requirements is treated in some technical detail. This analysis is then used as a basis for formulating a recommended position regarding use of X.25 in military networks.

2. HISTORY

In August of 1964, the RAND Corporation published an eleven volume series entitled, "On Distributed Communications," whose principle author was Paul Baran [1]. These volumes documented research sponsored by the US Air Force, which was motivated by work reported in a RAND paper [2] published in the summer of 1961. The objective of the research was to examine "the use of redundancy as one means of building communications to withstand heavy enemy attacks" [3]. The networking technique proposed for dealing with this difficult military communications problem was called a "Distributed Adaptive Message Block Network" by Baran, the "packet switching" terminology not yet having been coined. The concept exploited emerging new technology and included the assumptions that digital transmission and processing techniques would be used based upon anticipated dramatic reductions in the cost, size and power consumption of digital computer circuits by the 1970's. Baran's concept exploited these trends in technology to provide a network which would be richly connected and through which a path could be found to move traffic a "packet" at a time in store-and-forward fashion by a very flexible routing scheme if any path whatsoever existed between the origin and destination in an otherwise heavily damaged network. It should be

noted that the Baran oroposal was intended to carry all types of traffic, including voice and data, although most recent exploitations of packet switching have used the technique primarily for intercomputer and data terminal traffic.

In the late 1960's, the Defense Advanced Research Projects Agency (DARPA) undertook a program of research with the objective of providing a flexible method of communication among computers and between computers and terminals of different manufacture among their contractors who needed to collaborate on a number of research projects [4]. Connecting these facilities by an "ARPANET" would allow terminals access to a variety of hardware and software subsystems on a time-sharing basis and would support resource sharing among host computers. Bolt Baranek and Newman, Inc., was chosen as the prime contractor to develop the ARPANET and the first flow of traffic via the ARPANET in the packet-switched mode occurred in December of 1969. At that time the ARPANET was a small network consisting of four interconnected nodal processors, called IMPs, located at the Stanford Research Institute, the University of Utah and the University of California at Los Angeles and Santa Barbara. Since those early days, the ARPANET has undergone an evolutionary sequence of expansions and upgrades such that the network now stretches from the United Kingdom to Hawaii and consists of some one hundred nodes and three hundred host computer systems. The ARPANET has served both operational users who need its communications services and a research community which has continued to develop packet-switching technology for use by the military under the leadership of DARPA. Some of this research has involved refinement of original concepts (e.g., demonstrating the feasibility of packetized voice and development of improved routing algorithms), while some of it has involved striking out in new directions (e.g., development of broadcast techniques for local area, packet-radio and packet-satellite networks and development of a strategy for internetting diverse networks into a single interoperable community of users). Being a DoD Agency, DARPA's communications research objectives have been and continue to be strongly motivated by fundamental military communication requirements.

The success of the ARPANET experiments in packet-switching technology, together with complementary research at the National Physical Laboratory in England, provided motivation for a number of common carriers and other communications providers to consider offering packet-switched communication services on a public data network basis. Examples of early networks in this category include Telenet and Tymnet in the United States, Datapac in Canada, RCP and Transpac in France and an Experimental Packet-Switched Service (EPSS) in England. Today nearly every PTT and common carrier has operational or is developing a packet-switching communication service. These public data network services invariably are offered to the public on the basis of the concept of "virtual circuit." Even though different parts of a communications traffic stream may flow through a packet-switched network over different paths, the user can think in terms of traffic flow over a single "virtual" circuit connecting him to the host or terminal he is communicating with. It was natural that the concept of a virtual circuit, which first emerged from the ARPANET work, was adopted by the carriers as a way to describe and implement their service to prospective users, who were accustomed to thinking in terms of real circuits connecting them to the other members of their community of interest.

The virtual-circuit concept has undergone a further refinement by the development of protocols which require the establishment of a virtual circuit between users before any traffic can flow. This procedure leads to efficiencies (less overhead) in the handling of some types of traffic; it is unnecessarily restrictive for other types, however. Thus, while the concept of virtual circuit is useful, it should not become the only way to access a packet-switched service, which is sufficiently general to serve communications needs not well described in terms of circuits (e.g., broadcast services and especially, transaction services). For this reason, the concept of "connectionless communications" has been developed to convey the notion of a very general flow of traffic elements (packets or "datagrams") from a source to a destination (or set of destinations) over whatever paths may be available, different paths being used by different datagrams at different instants of time.

In the ARPANET, the host/IMP interface accepted independent "messages" from hosts, broke them into a sequence of packets, if necessary, routed them through the network, packet by packet, reassembled them at the destination IMP, and delivered the resulting messages in sequence relative to other messages sent by the same host to the destination. The host/IMP interface procedures, however, were datagram-like. Furthermore, to support packet voice traffic or other real-time applications, a nonsequencing datagram type of service on ARPANET was developed in 1974-75.

While the common carriers were in the process of adapting the packet-switching concept to the needs of public data network users, military planners continued to work towards the application of packet switching in their environment. In the United States, the DoD embarked on a study of military requirements for packet switching to determine whether a DoD common-user packet-switched service ought to be provided. In December of 1974 the DoD Data Internet Study report [5] provided the following recommendations: (1) that the ARPANET ought to be managed by the Defense Communications Agency (DCA) as an operational (though unsecure) communications network, and (2) that requirements for military communications justified the

development of a militarized (secure, survivable, interoperable, ...) packet-switched network modeled after the ARPANET. These recommendations were adopted and the ARPANET has been managed by DCA as an operational network since July of 1975. Moreover, several militarized packet-switched networks were developed and some are currently in operation (e.g., the World Wide Military Command and Control System (WWMCCS) Intercomputer Network (WIN), and the Community On-line Intelligence Network System (COINS)). However, the development of a fully militarized common-user packet-switched network for the DoD has turned out to be an elusive goal, the satisfaction of security and survivability objectives at reasonable cost being particularly difficult. The current DoD program for achieving this goal is called the Defense Data Network (DDN) Program. The DDN is being implemented in an evolutionary manner by incorporating the current ARPANET and several special-purpose ARPANET replicas into a single network, adding security and precedence features and an internetting capability.

3. REQUIREMENTS

Modern warfare places special demands on military communications that did not exist even a few years ago. These demands are related to the extreme degree of automation that has been or is being incorporated into military operations at all levels. In the tactical arena, all operations are being automated, from the aiming of weapons to the collection and processing of intelligence. Microcomputers are beginning to be used by the troops in the field [6], even as large main frames are used for strategic command and control, surveillance, warning and intelligence operations. It is not that the human is being replaced by machines (except for the most menial of functions); rather, human capability is being greatly extended by the leverage of automation. This trend is deliberate and is an important element in the US strategy of maximizing the military force achievable with limited manpower [7]. This trend has many analogies in other phases of modern society. Present day air travel simply would not be possible without computers to handle reservations and to control air traffic. We have become utterly dependent on automation, and this dependence is even stronger in military operations than in most other aspects of modern society. Moreover, effective computer communications greatly enhances the effectiveness of automated military operations; we are becoming extremely dependent upon these communications also. There is no turning back.

The modern military commander must not be deprived of automation and communications in the heat of battle. While some degree of manual back-up is necessary and even desirable, it is fundamental that the forces operate best when the capabilities at their disposal are the ones they are familiar with through training and exercises. Thus, we are faced with a bothersome fact of military life: the military requirement for communication of data to humans and machines demands that the communications operate at their best precisely when conditions are worst. They must provide effective communications to users on the move. They must operate when traffic demands far exceed the norms, when enemy action may have destroyed some facilities (which ones cannot be known in advance, of course) and when electronic countermeasures and sabotage are used to attack the system. They must communicate information that could be of great value to an enemy if intercepted and read. Communications management data must not only be protected from intercept (for it can be of considerable intelligence value to an enemy), it must also be protected from "spoofing" by an enemy who wishes to disrupt communications by interjecting false information. Moreover, these concerns apply not merely to tactical systems but to their strategic counterparts as well, since they are not immune from attack either.

In a word, the stress environment in which modern military communications must be able to operate is far different from the environment which has been assumed by national and international standards-setting committees. It would be folly, therefore, for military communicators to accept, without extremely careful consideration, standards which were not developed for the purpose of supporting the military mission.

Let us examine these special military requirements in a little more detail. First, a military system has a unique requirement for survivability. This requirement is, of course, closely related to the need for reliability. But the military system must survive in the face of hostile enemy action whose intent is to disrupt communications. There is also danger from collateral damage due to hostile military operations intended to attack nearby non-communications targets. To deal with this problem it is first necessary to define the likely threat (electronic countermeasures, direct attack, sabotage, etc.). Since one cannot predict the threat with a high degree of confidence, it is necessary to undertake corrective measures contrived to cover a wide variety of possible contingencies. For instance, architectures which are highly redundant and self-healing (or at least easily repaired) should be used for communications among the most important users. A variety of transmission methods may be used (e.g., satellite, wire line, line-of-sight microwave radio) in the hope that if one or more means is jammed or destroyed, one of the other means will still be available. The provision of survivable packet-switched communications under a variety of tactical and strategic conditions is still evolving. However, certain trends can be discerned. For instance, the use of richly connected, highly dispersed networks which continue to operate despite the partial loss of facilities is clearly indicated. The use of broadcast systems where possible seems very useful, because often the message can get through if any one of many receivers within

broadcast range can hear it. The use of many gateway interconnect points between different networks seems indicated.

Next let us consider the special military requirements of security. These are somewhat related to privacy requirements for public data services, but they are much more stringent. While there is considerable concern among carriers for protection of the privacy of medical records, for example, and for the use of methods which will counteract computer crime, for another example, the case for security in military systems is much more solid. So far, little money has been invested in the private sector towards the securing of computer communications. In contrast, considerable investment of defense R&D funds has been made in search of a solution to the security problem. Because of these trends, it must be admitted that, from the perspective of the military communicator, insufficient thought was devoted to the issue of security during the development of models by the Organization for International Standards (ISO) and the CCITT. This is not meant as a criticism of prevailing models or standards -- the military has not resolved this issue yet either -- it is a very difficult problem. But experience has shown that the application of a security mechanism on top of an existing architecture usually produces inferior results. One must really include security requirements in the architectural considerations from the start. Accordingly, military communicators should prudently conserve the option to devise a security architecture unencumbered by conformity to prevailing standards developed without sufficient consideration of military security requirements.

The reason the security requirement is so difficult to meet is that it is so broad and all encompassing. Not only must the enemy be deprived of reading the communicated text, he also must be denied access to communications management and control information (addresses, routing data, traffic statistics, etc.). Further, while access to the system must be simple for its intended users, it must be prohibited to interlopers who would be pleased to flood the system with meaningless data or otherwise interfere with its proper operation. The network must be able to deal with traffic of differing classification levels and categories, while guaranteeing that no traffic will be delivered to a user who has no access right to it. All of these security requirements must be satisfied without placing unreasonable operational constraints on the users who are trying to communicate in the midst of a military crisis.

Let us now turn to the requirement of dealing with the unusually high volume of traffic that occurs at the time of a crisis. This requirement has its counterpart in the private sector also; witness the difficulty experienced in communications during the Air Florida plane crash crisis in Washington, DC on January 13, 1982. But once again, military requirements are more stringent than those of the private sector. The value of communications service is at its maximum precisely at the time that the competing traffic load is maximum. If the military were to size its network for worst case crisis, the cost would be prohibitive and most of the facilities would lie idle most of the time. A solution to this dilemma has been employed in military communications for years; namely, the use of a precedence system which allocates communications resources to the most important uses, and requires less important traffic to receive minimal service until the crisis is over. In fact, well established criteria are invoked in today's military communications systems (e.g., AUTOVON, AUTODIN), and the method does work. In the case of packet switching, information is available in headers which can be more effectively used to discriminate in favor of high precedence traffic. While public packet-switched systems do provide for priority rationing of resources in times of crisis by setting aside optional format fields, no known existing implementation utilizes this optional feature at the present time.

Interoperability is an attribute greatly to be desired by both commercial and military systems. One of the major benefits of adhering to agreed upon standards is that this assures interoperability among the users of the standards. For this reason it is essential that military communications systems be compatible with commercial standards (such as CCITT X.25) which will enable them to interoperate with public packet-switched networks. In fact recent presidential directives require this type of interoperability. But the military cannot stop there. Requirements exist for interoperability among strategic and tactical networks, among military and commercial networks, and among national and allied networks (including NATO). Moreover, military long-haul terrestrial networks must be able to interoperate with local area networks, with tactical packet-radio networks and with packet-satellite networks of various types. Many of these other networks are based on broadcast protocols rather than circuit connection protocols. Examples are Ethernet, the DARPA Atlantic Packet Satellite Network (SATNET), tactical packet radio networks, Mitrebus, the Position Location Reporting System (PLRS) and the Joint Tactical Information Distribution System (JTIDS). While it is possible to interoperate with these broadcast media using a circuit-like protocol, it is awkward and inefficient to do so. Thus, exclusive use of virtual-circuit protocols fails to utilize inherent capabilities of these broadcast media which have been acquired at considerable effort.

The US DoD is moving in the direction of a multinetwork or "internet" architecture based on the concept of internet datagrams and gateway interconnections among diverse packet networks [8]. This issue is discussed in greater detail in a companion paper at this Symposium by V. Cerf [9]. The basic service provided by the

internet system is connectionless in nature. Virtual circuit services are provided on an end-to-end basis above the connectionless service through a transport layer Transmission Control Protocol (TCP) [10]. OARPA has supported extensive research and experimentation with the internet architecture, and has demonstrated its efficacy for interoperability of the high degree required to satisfy military requirements.

Based on concepts developed in the ARPANET [4], the Organization for International Standards (ISO) has established a model of protocol architectures based on a seven-layer structure [11]. This model, known as the Open System Interconnection (OSI) model, is shown in Figure 1. To properly deal with the emerging DoD Internet Architecture, it seems necessary to add an eighth "internet layer" to the ISO seven-layer model. This layer would be placed between the network layer and the transport layer. The DoD has already adopted a standard Internet Protocol (IP) [12,13] for this purpose and the US National Bureau of Standards has introduced it for consideration by ISO.

With this background of military requirements, the dilemma cited in the introduction can now be defined more explicitly: Should DoD reject the ISO seven-layer model in favor of its own model which supports its already existing and standardized protocols (IP and TCP)? The argument for doing this includes mention that IP and TCP have been implemented on a wide variety of computers in a wide variety of system software environments, with very extensive and favorable experience. [See Tables 1 and 2.] Or should the DoD adopt the ISO seven-layer model and attempt to have IP and TCP (or versions of them) adopted as "sublayer protocol standards" by some international standards organization? Or should the DoD abandon IP and TCP in favor of something else?

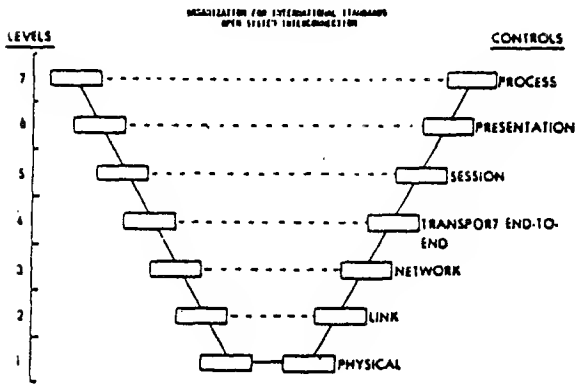

Figure 1. Protocol Layering

NETWORKS USING INTERNET ARCHITECTURE

ARPANET, WIN, DODIIS, COINS, MINET, EDN, BBNDIV6, DDN

PACKET RADIO (SRI, FT. BRAGG, SAC)

PACKET SATELLITE (SATNET, WBNET, MATNET)

LOCAL NETS (UNGERMANN-BASS NET/ONE, CHAOSNET, ETHERNET, MITREBUS, HYPERCHANNEL, PROTEON PRONET, BBN FIBERNET, SRINET, LEXNET)

CSNET (TELENET/ARPANET)

RSRE PILOT PACKET SWITCHING NET

UCL IPSS, PSS

DFVLR DATEX-P

USPS INTELPOST

TABLE - 1

INTERNET PROTOCOLS - IMPLEMENTATION

BBN

- DEC PDP-10/20 TENEX, TOPS-20
- DEC VAX UNIX (BERKELEY)
- DEC PDP-11/LSI-11 UNIX V7
- HONEYWELL H316 (ARPANET TAC)
- BBNCC C/30, C/70 UNIX, CMOS, TAC
- MOTOROLA MC68000 MINI-TAC
- HP HP3000 MPE

DEC

- DEC PDP-10/20 TOPS-20

COMSAT

- DEC LSI-11 RT-11, RSX-11, DCNET

3COM

- DEC PDP-11 UNIX
- DEC VAX UNIX
- ZILOG Z-8000 ZENIX

OTI

- DEC PDP-11 UNIX (NET FRONT END)
- DEC VAX VMS

SRI

- DEC LSI-11 MOS

RSRE

- DEC PDP-11 CORAL-66/VMOS

3RIVERS

PERQ

UCLA

IBM 360/91 OS/MVT
IBM 3033 OS/MVS

UNIV. MO

UNIVAC 1110 EXEC-8

MIT

HONEYWELL 6180 MULTICS (HONEYWELL TO MAINTAIN)
XEROX ALTO
IBM PERSONAL COMPUTER

FOONLY

FOONLY F-2, F-5 FOONEX

UNIV. WISCONSIN

IBM 4341

MITRE

OEC PDP-11 UNIX
ZILOG Z8000 CMOS (FRONT-END)

LOCKHEED

UNIVAC 1100 EXEC-8

TABLE - 2.

4. DOD PROTOCOL STANDARDIZATION POLICY

Department of Defense Instruction 4120.20 (DoDI 4120.20) entitled "Development and Use of Non-Government Specifications and Standards" [14] sets forth the prevailing policy of the DoD concerning adherence to national and international specifications and standards. This instruction implements Office of Management and Budget (OMB) Circular A-119 within the Department of Defense. As applied to protocol standardization, DoDI 4120.20 requires that national and international protocol standards be adopted as DoD standards in lieu of the development and promulgation of new documents. But the instruction also allows exceptions as necessary to provide for unique military requirements. In a memorandum dated 23 March 1982, Dr. Richard DeLauer, the Undersecretary of Defense for Research and Engineering, provided clarification of the DoD policy as it applies to protocol standards. While reiterating the need to utilize existing national and international standards where possible, he also reaffirmed the current policy of conformance to the existing DoD IP and TCP standards because "military requirements for interoperability, security, reliability and survivability are sufficiently pressing to have justified the development and adoption of TCP and IP in the absence of satisfactory non-government standards" [15]. The Defense Communications Agency, as DoD Executive Agent for Protocol Standardization, has been established as the authority to determine when military requirements justify the development and adoption of unique DoD protocol standards after making every effort to use existing standards. Moreover, the Executive Agent is to strive to inject DoD requirements into the national and international standards development process through participation in voluntary standards forums and through coordination with other US Government members of such forums (e.g., the National Communications System (NCS) and the National Bureau of Standards (NBS)). This influence is to be exerted with the objectives both of avoiding the need to develop and adopt unique DoD standards and of enabling their eventual replacement. (Further discussion of the US DoD Protocol Standards Program is provided by P. Selvaggi in a companion paper at this Symposium [16].)

5. AN ILLUSTRATIVE EXAMPLE: USE OF CCITT X.25 IN MILITARY NETWORKS

In this Section, we apply the litmus test of military requirements to a specific example of the fundamental dilemma. Let us examine, then, in some technical detail, the issue of applicability of CCITT recommendation X.25 to military communications networks.

The CCITT has developed a recommendation for the interfacing of subscriber computers to public packet-switched data networks. This recommendation is designated X.25 and specifies the procedures and formats by which subscriber data termination equipment (DTE) can exchange packets with the public data network data circuit termination equipment (DCE).

Recommendation X.25 also makes reference to lower level line control procedures and electrical interfacing options compatible with the so-called X.25 "packet level" interface protocol. These recommendations, taken together, constitute the body of the X.25 electrical, link and packet layer protocols which form the lowest three levels of the International Standards Organization's Open Systems Interconnection model portrayed in Figure 1. In addition to the X.25 recommendation, CCITT has also proposed other recommendations for interconnecting public data networks (X.75) and for interfacing computer terminals to public data networks (X.28, X.29).

The principal mode of operation of the X.25 interface is "virtual circuit" oriented. The subscriber DTE initiates a virtual circuit set-up procedure within the public data network

by sending to the public data network a connection request packet. Once the virtual circuit is established, the public data network returns a connection accept packet and the source and destination DTE's can then exchange data packets. A receiving DTE, upon receiving a connection request packet can accept or reject it. If the connection is rejected, the source DTE is informed of this by the public data network. The X.25 interface procedures include support for flow control between the public data network and the DTE, over and above the flow control exercised at the link level through HDLC procedures.

Insofar as the user is concerned, a virtual circuit allows him to reduce the header overhead in a multiple-packet message by referring to a logical channel group number and logical channel number that identify a call request packet in which call-specific parameters, such as addresses and facilities, are contained. A public data network which undertakes to offer the X.25 virtual circuit service must assure that packets are delivered in the same order in which they were received from the source DTE. Further, no packets may be lost or duplicated at the destination. If, as a consequence of some error or malfunction, the network cannot assure the integrity of the packet sequence, the virtual circuit is reset by sending a connection reset packet to both source and destination DTEs which must recover by setting up a new virtual circuit.

There exists a specification in the X.25 recommendation for a "datagram" mode of operation but it has not been implemented by any of the public data carriers and is not a mandatory service.

The X.75 protocol standard establishes formats and procedures which are used to interconnect public data networks. As a practical model, X.75 is merely an extension of the X.25 virtual circuit procedures which include, inter alia, provision for inter-network billing and grade of service indications. The X.75 procedures do not incorporate provisions for datagram mode of operation, nor do they specifically deal with internet routing procedures.

In summary, X.25 is a specification of a circuit-like interface for point-to-point communication between a pair of DTE's. It implies that the communication subnetwork will maintain the integrity and sequence of all packets sent into it by the source DTE. While this class of service is advantageous for some applications, it is not desirable for all applications, and it constrains the networking technology options so much that many important, existing military broadcast or semi-broadcast networks could not be used effectively, if at all, if the only interface to them had to adhere to the X.25 recommendations.

Examples of broadcast or semi-broadcast, military packet networks have been cited above. What is common about these networks is the fact that access to the common communication resource is shared in time, often by means of contention detection and resolution methods. This method of sharing access to a common communication capacity (e.g., radio channel, satellite transponder, coaxial cable) is extremely efficient for large numbers of bursty traffic sources, particularly mobile ones. By comparison, dedicated circuit-like sharing of these resources would be very wasteful of the capacity. We conjecture that in the future the private sector will share the military requirement for handling large numbers of bursty traffic sources efficiently, and that the connectionless protocols which are essential in the military environment will ultimately find wide acceptance commercially.

More to the point, however, is the fact that sequencing and integrity are services which are not desirable to build into such multiaccess networks. In order to sequence and maintain the integrity of the packets emitted by a source DTE, the subnetwork must be prepared to retransmit packets internally and to buffer them at least at the destination DCE to assure re-ordering, if necessary, before delivery to the destination DTE.

In mobile networks, or in systems where local jamming or other hostile action is likely, such services introduce congestion if a packet which has arrived at the destination cannot be delivered because it is not the "next" one. Sequencing also causes large variation in packet inter-arrival time at the destination DTE. The attempt to maintain integrity (i.e., not discard traffic) may also congest the network other than at the destination if the destination is out of contact (e.g., jammed, destroyed, beyond line-of-site, etc.) but this fact isn't yet known to the rest of the network.

The consequences of maintaining packet sequentiality within a store-and-forward network are increased total delay and increased delay variation. These effects interfere with the support of functions requiring real-time data services such as fire control and target tracking. Such real-time services need low delay, but not necessarily 100% integrity (i.e., some packet loss at the receiving DTE is acceptable). Furthermore, it has been found that the delay variations introduced by attempting to maintain sequencing and integrity within such networks make it impossible to support integrated packet voice services as part of the data network. The use of only X.25 virtual circuit service would negate a significant DoD investment in securable, low data-rate packet speech technology for integrated voice/data packet networks. Yet, current architectures for military tactical and

strategic communications call for at least limited use of packetized voice [17,18].

In addition to the unsatisfactory effects of its imposing sequencing and integrity requirements within multipoint communication networks, recommendation X.25 also does not deal explicitly with communication security and precedence, both of which are important to military communications.

Based on these considerations, many of the experimental US Department of Defense packet networks have been organized around the concept of datagrams. A datagram is a finite string of bits containing a header which typically indicates the destination address to which the bit string is to be sent, often indicating the source and other relevant information such as length, tyoe of transmission service desired, orecedence and so on. Datagrams typically are transoorted independently of each other by the packet networks, imposing few network mechanisms to imolement or use the simple and not necessarily reliable or sequential datagram service.

For the reasons given above, it does appear that military communication requirements would be well served, if all packet communication systems had to provide service only through interfaces meeting recommendation X.25 provisions. By the same token, however, it is essential that military communication systems be able to interconnect with and use networks offering only X.25 services.

During its exploration and development of packet-switching techniques, the US Department of Oefense has pursued the development and test of a layered protocol architecture which permits a broad range of different packet network types to be interconnected and used end-to-end, despite the very significant variations in the classes of services they offer, their different interfaces, speed of operation, throughput and maximum oacket sizes.

At the heart of this layered architecture is an Internet Protocol (IP) [12] which relies only on obtaining the most primitive datagram service from each constituent network of the Internet System. Figure 2 illustrates the relationships among the protocol layers of the DoD Internet Architecture.

The important differences between the DoD Internet Architecture and the networking models developed by CCITT and ISO are:

D1. The specific existence of an internet sublayer.

D2. The concept of gateways external to the communication subnetwork.

D3. The use of encapsulation to transport internet packets through intermediate networks.

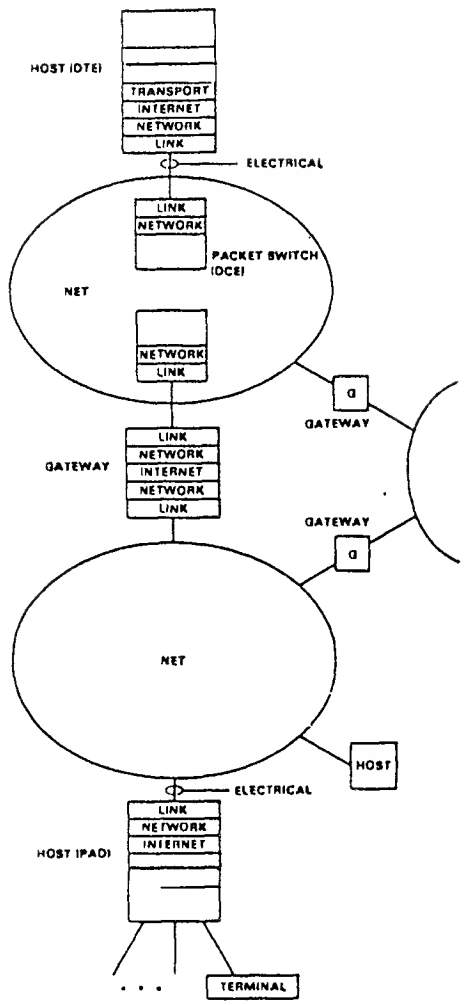

DOD INTERNET ARCHITECTURE
FIGURE 2

D4. The concept of gateway fragmentation and host (DTE) reassembly.

D5. The assumption that the basic network service is datagram and not virtual circuit.

D6. The provision for many network interfaces in addition to X.25.

D7. The explicit provision for security and precedence in the internet protocol sublayer.

While only the Internet Protocol and Transmission Control Protocol are ratified DoD standards, the other protocols are in widespread use in the experimental DoD Internet System which includes public networks (e.g., Telenet in the US and PSS and IPSS in the UK) as well as

experimental defense networks in the UK (Royal Signals and Radar Establishment (RSRE) Pilot Packet Switched Network (PPSN)) and Norway (Norwegian Defense Research Establishment -- NORENET).

Creation of the separate gateway system has made it possible to incorporate into the DoD Architecture mechanisms for recovering from partitioning of a communication subnetwork by routing traffic through the internet gateway system as illustrated in Figure 3 [9].

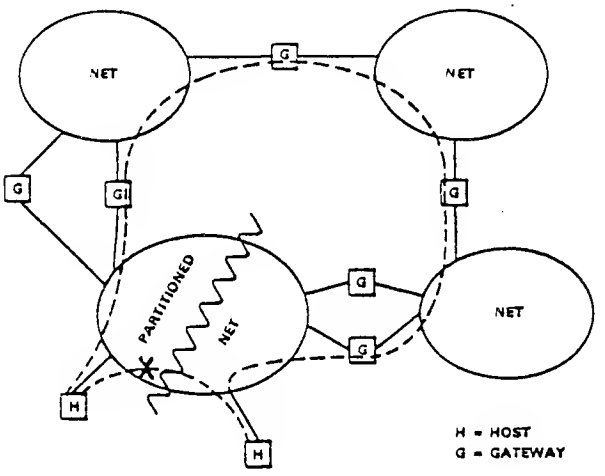

PARTITIONED NET RECOVERY
FIGURE 3

The encapsulation and fragmentation mechanisms at the internet level have allowed internet routing to be decoupled from the problem of accommodating varying maximum packet sizes in each network. The strategy allows the network packet size to be optimized to the particular switching and transmission technology (as well as local communication and propagation conditions) in each network (Figure 4).

By assuming only datagram services from each constituent network, the DoD Internet Architecture is able to support a broader range of applications including real-time packet voice. Each internet packet can include an indication of the type of service it needs, which might trigger the use of virtual circuits on some networks which provide the service, but in general, virtual-circuit-like service is provided by the Transmission Control Protocol at the transport layer, outside the collection of interconnected packet subnetworks.

The simplifying datagram service assumption also makes it easier to use multiple gateways to share traffic loads since it isn't necessary to maintain sequencing and integrity of any particular virtual circuit passing across a given network. Packets can be switched to alternate gateways as appropriate to share their capacity. This also helps to speed up recovery when gateways fail without necessarily requiring action on the part of the source DTE (host).

While details of the security architecture for the DoD Internet System are classified, it is nevertheless necessary for the military to work within standardization forums to achieve basic architectures which can serve both the military need for security and the needs of the private sector for privacy and integrity. It should be possible to define a generic architecture which serves the needs of the military and nonmilitary users, and which military communicators can adapt to the needs of military security.

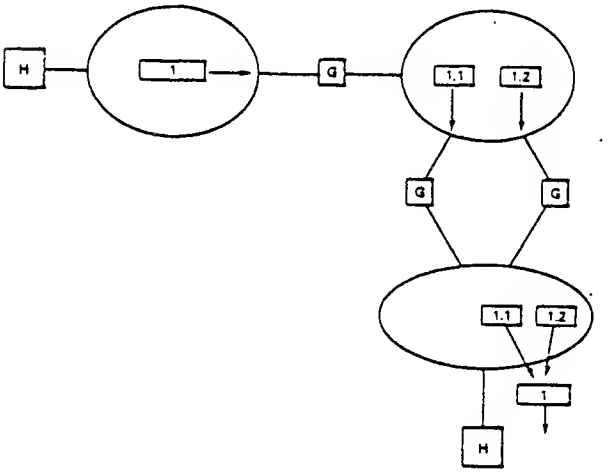

INTERNET PACKET FRAGMENTATION AND REASSEMBLY
FIGURE 4

The US Defense Advanced Research Projects Agency (DARPA) is currently conducting experiments with its counterparts in Norway (Norwegian Oefense Research Establishment), the United Kingdom (Royal Signals and Radar Establishment), and Germany (OFVLR -- the German Air and Space Research Agency) on the use and further development of the Internet Architecture. These experiments are relevant to the on-going discussion among NATO countries concerning NATO standards for packet communication, and their results should be factored into any decisions and agreements reached within the NATO normalization and standardization process.

In view of the foregoing, an acceptable US or NATO standard network architecture must be able to make use of, but not be limited to, networks providing interfaces meeting the CCITT Recommendation X.25. In particular, provision for nonvirtual circuit modes of operation are considered mandatory to support transaction or real-time applications in an internetwork environment.

6. RECOMMENDED STRATEGY

We have attempted to show that current circumstances compel military networking systems to utilize and supoort certain unique protocols in addition to prevailing national and international standard protocols in order to meet their special requirements. This is true, at least as far as IP and TCP protocols are concerned, despite the obvious disadvantages of the use of unique standards.

The reason the military finds itself in a dilemma at the present time is that military requirements have been inadequately represented in the standards forums that have formulated currently prevailing standards. Thus, if the military finds itself wishing that the ISO model for Open Systems Interconnection incorporated an internet protocol layer, it is because they have not convinced ISO and CCITT of this requirement. Similarly, if there is a sufficient need for a connectionless (datagram) protocol to parallel the CCITT X.25 virtual circuit orotocol, then military users and others who share that need ought to be able to convince the international standards setting bodies of that requirement. The demonstration of sufficient need requires considerable justification, of course. One can argue that many of the military requirements not well served by prevailing standards have their counterparts in the commercial sector -- often a perceived military need is the precursor of a later recognized nonmilitary need. Moreover, every nation has similar military communication needs and most PTT's count their military users among their imoortant customers. But these arguments have never been accepted in international standardization forums. Perhaps it can be said that military planners did not take the movement toward data communications standards seriously enough. In any case, a prescription for future work is needed.

The authors recommend the following strategy for the future of data communications protocol standardization for military communications planners:

a. The US DoD should continue to maintain the Internet Protocol (IP) and the Transmission Control Protocol (TCP) standards for the present because of their superior ability to satisfy established military requirements.

b. Other national defense departments and ministries (including NATO) should consider adopting these protocol standards for the same reasons.

c. Other standards, such as CCITT X.25, should also be adopted for use where appropriate within the DoD and other military establishments, and brought under configuration control, so that the military users will be able to interface to public networks when that is required to satisfy interoperability and service restoral requirements as stipulated in US national policy.

d. In particular, a datagram counterpart to the current X.25 virtual circuit standard ought to be defined to serve those (including the military) who need connectionless services. (A datagram service has already been defined as an option under X.25, but it is not well suited to the needs of the military or any other users, so will tend to remain unused unless significantly modified.)

e. As a means of insuring interoperability, configuration control of protocols and their implementations needs to be assigned to an authority so that means will be available to maintain updated authoritative records and to certify that proposed implementations are correct.

f. Representatives of Allied military establishments ought to work within the existing national and international standardization forums to foster understanding of their requirements and support for their needs for protocol standards.

g. Defense R&D activities should prepare an architecture for protocols that adopts the terminology and structure of the ISO seven-layer model insofar as this is feasible-while still satisfying valid military requirements. This architecture should then be publicized in order to influence the evolution of national and international standards to the extent possible.

h. Defense R&D activities should pursue the development of new protocols in collaboration with their commercial counterparts, and within the framework of prevailing national and international standards, such that future protocol standards will tend to converge toward accommodation of military and commercial requirements.

i. Defense officials should work closely with their civilian government counterparts (in the US, the DoD should work with the National Communication System and the National Bureau of Standards) to ensure that military needs are taken into consideration as a subset of total government needs for standards.

7. SUMMARY

This paper explores a basic dilemma that faces the military communicator: Should military data communications systems use special protocol

standards unique to the military, or should they use prevailing national and international standards? The latter alternative is to be preferred, provided the military mission can be accomplished. We have argued that in the oresent environment it is not possible to satisfy military requirements without including some uniquely military protocol standards. We have also argued that this unfortunate situation ought to be a temporary one, with military and civilian protocol standards converging towards a single optimum set in the future. The objectives of US OoO research and development programs in this area, with its past history of significant contributions, are already being oriented toward achieving this convergence. There is some optimism that this convergence will, in fact, occur. In the first place, military needs tend to be precursors of future civilian needs: today's military security requirement is precursor to tomorrow's commercial need for privacy and data integrity; today's military need for broadcast and real-time services is precursor to tomorrow's commercial need for cellular systems and transaction processes. In the second place, pressure from industry on military communications planners and on common carriers for convergence on a single set of standards will be intense. The bottom-line conclusion of this paper is that, granted the temporary need for special military protocol standards, it is time for industry and the military to close ranks by making the desired convergence on a single set of protocol standards a firm requirement of the future efforts of both communities.

8. ACKNOWLEDGMENT

The authors wish to express their gratitude to many colleagues who reviewed the manuscript for this paper and contributed beneficial suggestions for improvement.

REFERENCES

[1] The RANO Corporation, "On Distributed Communications," Eleven Volumes, August 1964.

[2] The RAND Corporation, Paper P-2626, 1961.

[3] The RAND Corporation, "On Oistributed Communications," op cit, Vol. I, "Introduction to Oistributed Communications Networks," by Paul Baran (RM-3420-PR), p. iii.

[4] Defense Advanced Research Projects Agency, "A History of the ARPANET: The First Decade," 1 April 1981.

[5] Department of Defense, DoO Data Internet Study Phase II Report, Oecember 1974.

[6] Schneider, William P., "Small Computers in the Army: An Apple a Day to Keep the Soviets Away," Signal, February 1982, pp. 39-43.

[7] Perry, William J. and Roberts, Cynthia A., "Winning Through Sophistication: How to Meet the Soviet Military Challenge," Technology Review, July 1982, pp. 26-35.

[8] Davies, B. H., and Bates, A. S., "Internetworking in the Military Environment," INFOCOM 82, IEEE Press, 1982, pp. 19-29.

[9] Cerf, Vinton G., "The US Department of Defense Internet Architecture," SHAPE Technical Center Symposium on Interoperability of Automated Data Systems, The Hague, November 1982.

[10] Defense Advanced Research Projects Agency, "Transmission Control Protocol," RFC 793, September 1981.

[11] International Standards Organization, "Reference Model of Open System Interconnection," ISO/TC97/SC16/N227, 1979.

[12] Defense Advanced Research Projects Agency, "Internet Protocol," RFC 791, September 1981.

[13] Defense Advanced Research Projects Agency, "Internet Control Message Protocol," RFC 792, September 1981.

[14] Department of Defense, DoDI 4120.20, "Oevelopment and Use of Non-Government Specifications and Standards," 28 December 1976.

[15] DeLauer, Richard O., "DoO Policy on Standardization of Host-to-Host Protocols for Oata Communications Networks," 25 March 1982.

[16] Selvaggi, Philip S., "The US Department of Defense Protocol Standards Program," SHAPE Technical Center Symposium on Interoperability of Automated Data Systems, the Hague, November 1982.

[17] Tactical Information Exchange (TIE) Working Group, "Tactical Information Exchange (TIE) Framework Development," edited by R&O Associates, ROA-TR-117100-001, October 1981.

[18] LaVean, Gilbert E. and Sonderegger, Ronald E., "A Communication System Architecture for Interoperable Systems," International Telemetering Conference," San Oiego, CA, 28-30 September 1982.